

### Abstract of the Disclosure

A method for reorganizing rows from a partitioned database table is disclosed. The partitioned database table includes a plurality of populated partitions. Rows are organized in each of the populated partitions in accordance with a first value associated with each row. A file context is created for each partition of a subset of the populated partitions, each file context stores at least location data for a row in the partition and the first value associated with the row. Rows are merged from the subset of partitions into a single first-merge partition in order of the first value associated with each row. The creation of file contexts and merging of rows in the corresponding partitions is repeated until the subsets have included all populated partitions.